Optimal Deterministic Algorithm Generation

Alexander Mitsos[a], Jaromił Najman[a], Ioannis G. Kevrekidis[b], *

[a]AVT Process Systems Engineering (SVT), RWTH Aachen University, Turmstrasse 46, Aachen, 52064, Germany, amitsos@alum.mit.edu ORCID: 0000-0003-0335-6566

[b]Department of Chemical and Biological Engineering and Engineering & Program in Applied and Computational Mathematics, Princeton University, A319 Engineering Quad;Princeton, NJ 08544, yannis@princeton.edu, Phone: 609-258-2818 (also IAS-TUM, Garching and Zuse Institut, FU Berlin, Germany.)

arXiv:1609.06917v1 [math.OC] 22 Sep 2016

A formulation for the automated generation of algorithms via mathematical programming (optimization) is proposed. The formulation is based on the concept of optimizing within a parameterized family of algorithms, or equivalently a family of functions describing the algorithmic steps. The optimization variables are the parameters - within this family of algorithms- that encode algorithm design: the computational steps of which the selected algorithms consists. The objective function of the optimization problem encodes the merit function of the algorithm, e.g., the computational cost (possibly also including a cost component for memory requirements) of the algorithm execution. The constraints of the optimization problem ensure convergence of the algorithm, i.e., solution of the problem at hand. The formulation is described prototypically for algorithms used in solving nonlinear equations and in performing unconstrained optimization; the parametrized algorithm family considered is that of monomials in function and derivative evaluation (including negative powers). A prototype implementation in GAMS is developed along with illustrative results demonstrating cases for which well-known algorithms are shown to be optimal. The formulation is a mixed-integer nonlinear program (MINLP). To overcome the multimodality arising from nonconvexity in the optimization problem, a combination of brute force and general-purpose deterministic global algorithms is employed to guarantee the optimality of the algorithm devised. We then discuss several directions towards which this methodology can be extended, their scope and limitations.

Optimization |Nonlinear Equations |Algorithms |Optimal Control

Computational engineering and science has realized major advances, in large part due to ingenious numerical algorithms. The development of algorithms is thus of crucial importance and requires substantial resources and time. Typically multiple algorithms exist for a given task. Comparisons of algorithms that tackle a given family of problems is sometimes performed theoretically and sometimes numerically. It would be desirable to automatically generate algorithms and have a guarantee that these are the best for a class of problems or for a given problem. We propose the use of numerical optimization to design *optimal algorithms*, i.e., algorithms that perform better than any conceivable alternative. More precisely, each iteration of the algorithms is interpreted, in an input-output sense, as a function evaluation. Then an optimization problem is formulated that finds the best algorithm among a family of algorithms/functions, where "best" is determined in a predefined metric.

There is large literature on automatic generation of algorithms but in most cases the idea of genetic algorithms is employed. Genetic algorithms first generate a set of algorithms evaluated on a test set built of elementary components of potentially promising algorithms and subsequently combine the most promising ones in order to obtain a new elitist algorithm [3, 15, 16]. The work by [15, 16] covers basic methods and ideas in the field of genetic programming. Automatic algorithm generation for retrieval, classification and other data manipulation has been proposed by [17] based on statistical data analysis and other methods. For instance, in [14] algorithms for the well-known knapsack problem are automatically generated by genetic algorithms. First algorithms for smaller knapsack instances are generated and then compared to algorithms computed on larger test sets. In [14] an increase in efficiency is noted if larger instantiations for the algorithm generation are used.

Numerical algorithms are ubiquitous in tackling scientific and technical tasks and their relevance is increasing rapidly. Even daily activities, such as finding the optimal route via a navigation system, employ algorithms. Typically algorithms are developed by creative individuals in incremental or revolutionary steps, thus requiring substantial effort and not guaranteeing that the algorithm is the most suitable for the task at hand. Herein an approach for automated algorithm generation is proposed. The design of the algorithm itself is formulated as an optimization problem which in turn is solved using numerical methods (algorithm-building algorithms) that guarantee optimality. Thus, the selected algorithms are the best possible in a predefined metric among a family of possible algorithms considered.

In [23] work is presented on automatic generation of parallel sorting algorithms. They combine known sorting algo-

Corresponding author: *Ioannis G. Kevrekidis

rithms to obtain a best performing algorithm for a certain input. The well known art gallery problem (AGP) is discussed by [29], where an iterative algorithm is developed and in that sense resembles our proposal. Moreover, there is substantial work on automatic code and algorithm generation for certain problems, e.g., [1, 2, 11, 24]. While this approach is fundamentally different, in a sense it resembles the proposal herein: both require an external algorithm in order to compute a sub-result. Many depend on algebraic operations regarding matrices, e.g., QR decomposition, diagonal transformation or eigenvalue computation. Many of those operations are well documented in the Thesis of [5]. Finally, the problem of tuning the parameters of existing algorithms through optimization has been considered [18]

Consider an iterative algorithm, like the time-honored Newton-Raphson for finding roots of algebraic equations, e.g., in one dimension $x_{n+1} = x_n + f(x_n)/f'(x_n) \equiv x_n + u_n(x_n)$. The iterates can thus be seen as a "control action" u_n that depends only on the current guess, i.e., "state-dependent feedback" and the same interpretation can be given to most algorithms. But if algorithms are feedback laws (towards a prescribed objective), then we do have celebrated mathematical tools for computing optimal feedback, e.g., HamiltonJacobiBellman. It would make sense to use these tools to devise optimal algorithms.

The key goal of the manuscript is to formulate a systematic process for obtaining optimal algorithms. To the best of our knowledge, there exists no other proposal in the literature to use deterministic global optimization for finding an algorithm that is optimal w.r.t. a certain objective. Optimality is determined based on a desired measurable performance property which is encoded as the objective of the optimization problem. For instance this can be the computational expense of the algorithm or the memory requirement. Upon convergence of our formulation, the resulting algorithm is optimal among a considered family of algorithms which in turn is encoded using the variables (or degrees of freedom) of the optimization problem. The optimization formulation is completed by the constraints which ensure that only feasible algorithms will be selected among the family of algorithms considered, i.e., algorithms that converge to a solution of the problem at hand. The formulations can be cast as nonlinear programs (NLP) or mixed-integer nonlinear programs (MINLP).

One of our proof of concept illustrations involves, e.g., devising optimal algorithms that perform local unconstrained minimization of scalar functions. The class of algorithms within which optimality is sought involve evaluations of the function and its first two derivatives; they can be composed of up to ten iterations of a monomial involving these quantities, what we call a monomial-type algorithm. The well-known steepest descent is shown to be optimal in finding a stationary point of a univariate fourth-order polynomial, whereas it is infeasible for the two-dimensional Rosenbrock objective function.

The remainder of the manuscript is structured starting with the simplest possible setting, namely solution of a single nonlinear equation by a particular class of algorithms. Two immediate extensions are considered to equations systems and to local unconstrained optimization. Subsequently, numerical results are presented along with some discussion. Potential extension to more general problems and algorithm classes are then discussed as well as limitations and relation to computational complexity. Finally, key conclusions are presented.

Definitions and Assumptions

For the sake of simplicity, first, solution of a single nonlinear equation is considered as the prototypical task an algorithm has to tackle.

Definition 1. Let $x \in X \subset \mathbb{R}$ and $f : X \to \mathbb{R}$. A solution of the equation is a point $x^* \in X$ with $f(x^*) = 0$. An approximate solution of the equation is a point $x^* \in X$ with $|f(x^*)| \leqslant \varepsilon$.

We will first make use of relatively simple algorithms:

Definition 2. Algorithms operate on the variable space $\mathbf{x} \in X \subset \mathbb{R}^{n_x}$ and start with an initial guess $\mathbf{x}^{(0)} \in X$. A given iteration it of a simple algorithm amounts to calculating the current iterate as a function of the previous $\mathbf{x}^{(it)} = g_{it}\left(\mathbf{x}^{(it-1)}\right)$, $g_{it} : X \to X$. Therefore, after iteration it, the algorithm has calculated $\mathbf{x}^{(it)} = g_{it}\left(g_{it-1}\left(\ldots(g_2(g_1(\mathbf{x}^{(0)})))\right)\right)$. A special case are algorithms that satisfy $g_{it_1}(x) = g_{it_2}(x)$ for all it_1, it_2, i.e., use the same function at each iteration.

In other words, algorithmic iterations can be seen as functions and algorithms as composite functions. This motivates that finding an optimal algorithm is an optimal control problem. In some sense the formulation finds the optimal function among a set of algorithms considered. One could thus talk of a "hyper-algorithm" implying a (parametrized) family of algorithms, or possibly the span of a "basis set" of algorithms that are used to identify an optimal algorithm. In the following we will assume $X = \mathbb{R}^{n_x}$ and consider approximate feasible algorithms:

Definition 3. An algorithm is *feasible* if it solves a problem; otherwise it is termed *infeasible*. More precisely: it is termed *feasible in the limit* for a given problem if it solves this problem as the number of iterations approaches ∞; it is termed *finitely feasible* if it solves the problem after a finite number of iterations; it is termed *approximate feasible* if after a finite number of iterations it solves a problem approximately. In all these cases the statements holds for some (potentially open) set of initial conditions. A feasible algorithm is *optimal* with respect to a metric if among all feasible algorithms it minimizes that metric.

Throughout the report it is assumed that $f : X \subset \mathbb{R} \to \mathbb{R}$ is sufficiently smooth in the sense that if algorithms are considered that use derivatives of a given order, then f is continuously differentiable to at least that order. The derivative of order j is denoted by $f^{(j)}$ with $f^{(0)} \equiv f$.

A key point in our approach is the selection of a class of algorithms: a (parametrizable) family of functions. It is possible, at least in principle, to directly consider the functions g_{it} and optimize in the corresponding space of functions. Alternatively, one can identify a specific family of algorithms, e.g., those based on gradients, which includes well-known algorithms such as Newton and explicit Euler:

Definition 4. Consider problems that involve a single variable $x \in X \subset \mathbb{R}$ and $f : X \to \mathbb{R}$. *Monomial-type algorithms* are those that consist of a monomial $g_{it}(x) = x + \alpha \Pi_{j=0}^{j_{max}} \left(f^{(j)}(x) \right)^{\nu_j}$ of derivatives of at most order j_{max}, allowing for negative powers.

Development

We first develop an optimization formulation that identifies optimal algorithms of a particular problem: the solution of a given nonlinear scalar equation, Definition 1. We will consider approximately feasible solutions and simple algorithms that use the same function at each iteration. If the optimization metric accounts for the computational cost of the algorithm, then it is expected that algorithms that are only feasible in the limit will not be optimal since their cost will be higher than those of finitely feasible algorithms.

Infinite Problem

We are interested in devising algorithms as the solution of an optimization problem. The key idea is to include the iterates of the algorithm $x^{(it)}$ as (intermediate) *optimization variables* that are determined by the algorithmic iteration. By imposing a maximal number of iterations, we have a finite number of variables. The algorithm itself is an optimization variable, albeit an infinite one (optimal control). Moreover, an end-point constraint ensures that only feasible algorithms are considered. Finally, an optimization metric is required which maps from the function space to $\mathbb{R}$, such as the computational cost. We will make the assumption that the objective is the sum of a cost at each iteration which in turn only depends on the iterate, and implicitly on the problem to be solved.

$$\begin{aligned} &\min_{g, x^{(it)}, it_{con}} \sum_{it=1}^{it_{con}} \phi \left(g \left(x^{(it)} \right) \right) \\ &\text{s.t. } x^{(it)} = g \left(x^{(it-1)} \right), \quad i = 1, 2, \ldots, it_{con} \\ &\left| f \left(x^{(it_{con})} \right) \right| \leqslant \epsilon \end{aligned} \tag{1}$$

it_{con} corresponds to the iteration at which the desired convergence criterion is met. The optimal point found embodies an optimal algorithm g^*; depending on the method used it may be global, approximately global or local optimum.

For notational simplicity, we have assumed that the minimum is attained. It is outside the scope of this manuscript to determine if the infimum is attained, e.g., if g can be described by a finite basis of a finite vector space.

The intermediate variables can be considered as part of the optimization variables along with the constraints or eliminated. Note the analogy to full discretization vs. late discretization vs. multiple shooting in dynamic optimization [21, 4] with known advantages and disadvantages.

(1) is not a regular optimal control problem as the number of variables is not a priori known. There are however several straightforward approaches to overcome this problem, such as the following. A maximal number of iterations it_{max} is considered a priori and dummy variables for $it_{con} < it < it_{max}$ are used along with binary variables capturing if the convergence has been met. The cost is accounted only for the iterates before convergence:

$$\min_{g, x^{(it)}} \sum_{it=1}^{it_{max}} yres^{(it)} \phi \left(g \left(x^{(it)} \right) \right)$$

$$\text{s.t. } x^{(it)} = g \left(x^{(it-1)} \right), \quad it = 1, 2, \ldots, it_{max} \tag{2}$$

$$y_{res}^{it} = \begin{cases} 0 & \text{if } \left| f \left(x^{(it)} \right) \right| < \varepsilon \\ 1 & \text{else} \end{cases} \tag{3}$$

$$y_{res}^{it_{max}} = 0 \tag{4}$$

where satisfaction of the last constraint ensures convergence. In the supporting information an alternative is given. In the formulation above each iteration uses the same function g. Allowing different function is of interest, e.g., to devise hybrid iteration algorithms, such as few steepest descent steps followed by Newton steps.

Finite Problem

We now use the family of monomial-type algorithms, Definition 4 as our ensemble of algorithms. In analogy to the infinite case we assume that the objective depends on the evaluation of derivatives (of the result of the algorithm with respect to its inputs), which in turn only depends on the iterate, and implicitly on the problem to be solved. To account for the expense of manipulating the derivatives, e.g., forming the inverse (recall that simple algorithms scale cubically in the size), we assume that the cost also involves an exponent:

$$\min_{\nu_j,\alpha,x^{(it)},it_{con}} \sum_{it=1}^{it_{con}} \sum_{j=0}^{j_{max}} \phi_j\left(f^{(j)}\left(x^{(it)}\right),\alpha,\nu_j\right)$$
$$\text{s.t. } x^{(it)} = x^{(it-1)} + \alpha\Pi_{j=0}^{j_{max}}\left(f^{(j)}(x^{(it-1)})\right)^{\nu_j}, it = 1,\ldots,it_{con}$$
$$\left|f\left(x^{(it_{con})}\right)\right| \leqslant \epsilon \qquad (5)$$

where again the number of variables is not known a priori, see above. Note that since the exponents are integer-valued we have a mixed-integer nonlinear program (MINLP) and thus, for a fixed f we expect that the minimum is attained. Again the intermediate variables can be considered as part of the optimization variables along with the constraints, or they can be eliminated. The optimal point found gives optimal stepsize α^* as well as exponents ν_j^*. Allowing for a different algorithm at each step ($\nu_j^{it_1} \neq \nu_j^{it_2}$), the number of combinations dramatically increases but there is no *conceptual* difficulty. The form is not common for an MINLP but can be easily converted to a standard one as shown in the supporting information.

It may not always be easy to estimate the cost $\phi_j\left(f^{(j)}\left(x^{(it)}\right),\alpha,\nu_j\right)$. For instance consider Newton's method augmented with a line-search method $x^{(it)} = x^{(it-1)} + \alpha^{it}\frac{f(x^{(it-1)})}{f^{(1)}(x^{(it-1)})}$. For each major iteration, multiple minor iterations are required, with evaluations of f. If the cost of the minor iterations is negligible, the computational cost is simply the evaluation of f, its derivative and its inversion. If we know the number of minor iterations then we can calculate the required number of evaluations. If the number of iterations is not known, we can possibly obtain it by the step size, but this requires knowledge of the line search strategy. In some sense this is an advantage of the proposed approach: the stepsize is not determined by a line search method but rather by the optimization formulation. However, a challenge is that we need to use a good cost for the stepsize; otherwise the optimizer can select an optimal α for that instance which is not necessarily sensible when the cost is accounted for. To mimic the way conventional algorithms work, in the numerical results we allow discrete values of the steplength $\alpha^{(it)} = \pm 2^{-\bar{\alpha}^{(it)}}$ and put a bigger cost for bigger $\bar{\alpha}$ since this corresponds to more evaluations of the line search.

The algorithm is fully determined by selecting ν_j and $\bar{\alpha}^{(it)}$. So in some sense the MINLP is a purely integer problem: the variables $x^{(it)}$ are uniquely determined by the equations. To give a sense of the problem complexity consider the setup used in the numerical case studies. Assume we allow derivatives of order $j \in \{0,1,2\}$ and exponents $\nu_j \in \{-2,-1,0,1,2\}$ and these are fixed for each iteration it. This gives $5^3 = 125$ basic algorithms. Further we decide for each algorithm it if the stepsize α is positive or negative, and allow $\bar{\alpha}^{(it)} \in \{0,1,\ldots,10\}$, different for each iteration it of the algorithm. Thus we have 2×10^5 combinations of stepsizes for each algorithm and 25×10^6 total number of combinations.

Equation Systems

A natural extension of solving a single equation is to solve a system of equations.

Definition 5. Let $\mathbf{x} \in X \subset \mathbb{R}^{n_x}$ and $\mathbf{f} : X \to \mathbb{R}^{n_x}$. A solution of the equation system is a point $\mathbf{x}^* \in X$ with $\mathbf{f}(\mathbf{x}^*) = \mathbf{0}$. An approximate solution of the equation is a point $\mathbf{x}^* \in X$ with $||\mathbf{f}\left(\mathbf{x}^*\right)|| \leqslant \varepsilon$.

In addition to more cumbersome notation, the optimization problems become much more challenging to solve due to increased number of variables and also due to the more expensive operations. Obviously the monomials need to be defined appropriately; for instance the inverse of the first derivative corresponds to the inverse of the Jacobian, assuming it has a unique solution. This in turn can be written as an implicit function or an equation system can be used in the optimization formulation. For the numerical results herein we use small dimensions (up to two herein) and analytic expressions for the inverse. In the supplementary material we discuss an approach more amenable to higher dimensions.

Finding Optimal Local Optimization Algorithms

It is relatively easy to adapt the formulation from equation solving to local solution of optimization problems. The only substantial change is to replace the end-point constraint with some desired termination criterion such as an approximate solution of the KKT conditions. For unconstrained problems the corresponding criterion is stationarity of the objective $\frac{\partial f}{\partial x_{ix}} = 0$. Similarly to equation systems increased dimensionality ($n_x > 1$) implies that vectors and matrices arise and operations need to be appropriately formulated.

Numerical Results

Method

We develop a prototype implementation in GAMS [9], inspired by [20]. We used GAMS 24.5 and tested both local and global methods, mostly KNITRO [10] and BARON [28].

As aforementioned we encode the choice of derivative as a set of binary variables. Two problems are considered, namely solution of equation systems and minimization, and in both cases we consider one-dimensional and two-dimensional problems. The minimization is assumed unconstrained but since we impose bounds on the variables we essentially assume knowledge of an open box in which optimal solutions are found. For two-dimensional problems,

for simplicity, we do not allow second derivatives. We allow 10 steps of the algorithm. We require that at least the last iterate meets the desired convergence criteria.

We tried several runs, mostly with short times (order of minute) but some with longer. The key findings is that KNITRO often fails to find feasible algorithms, even when they exist. BARON often also fails to find feasible algorithms within few minutes. In several cases BARON manages to find feasible algorithms and converge the lower bound but in many cases the lower bound stays at low levels, i.e., the certificate of optimality is relatively weak. Often the convergence behavior of BARON is improved if it is initialized with KNITRO. In some cases convergence of the optimization was improved when we enforced that residuals are decreasing with iteration $res^{it} < res^{it-1}$. It is noteworthy that in early computational experiments the optimizer also managed to furnish unexpected solutions that resulted in refining the formulation, e.g., choosing a stepsize α that resulted in direct convergence before discrete values of α were enforced.

Due to the relatively difficult convergence of the optimization and the relatively small number of distinct algorithms (125), we solved each problem for each algorithm in three subsequent runs: first KNITRO to find good initial guesses, then BARON without convergence criterion (i.e., without imposing $y_{res}^{it_{con}} = 0$) and then BARON with convergence criterion. Each of the runs was given 60 seconds in a first computational experiment and 600 seconds in a second. With the small CPU time some runs converged to the specified optimality tolerance of 0.1 (both absolute and relative), while most did not; with the higher CPU more algorithms were proven to be infeasible while, for some, feasible solutions were found. Convergence to the global optimum is not guaranteed for all cases. The main findings of this explicit enumeration is that most algorithms are infeasible. For equation-solving, in addition to Newton's algorithm some algorithms with second derivatives are feasible. During the loop some runs resulted in internal errors.

The algorithms for unconstrained minimization are instructive. In the problem $\min_{x\in[-2,2]} x^4 + x^3 - x^2 - 1$, with $x^{(0)} = 0.1$ the steepest descent algorithm is optimal. It is interesting to note that the algorithms do not furnish a global minimum: in the specific formulation used optimal algorithms are computationally cheap algorithms that give a stationary point. In contrast, for the minimization of the Rosenbrock function in 2d $\min_{\mathbf{x}\in[-2,2]^2} 100(x_2 - x_1^2)^2 + (1 - x_1)^2$, with $x_1^{(0)} = 0.7$, $x_2^{(0)} = 0.75$ steepest descent is not feasible within ten iterations which is well-known behavior. From the algorithms to solve a single equation it is interesting to note that several algorithms were furnished. Two particularly interesting algorithms identified (with $\alpha < 0$) are

$$x^{(it)} = x^{(it-1)} + \alpha \frac{f^{(0)}(x^{(it-1)}) f^{(1)}(x^{(it-1)})}{f^{(2)}(x^{(it-1)})}$$

$$x^{(it)} = x^{(it-1)} + \alpha \frac{f^{(0)}(x^{(it-1)}) \left(f^{(1)}(x^{(it-1)})\right)^2}{\left(f^{(2)}(x^{(it-1)})\right)^2}.$$

which can be seen as combination of Newton's algorithms for equation solving and unconstrained optimization. The steplengths chosen depend on iteration. It would be interesting to confirm optimality through a HJB approach.

Discussion of Possible Extensions

In principle, the proposed formulation can be applied to any problem and algorithm. In this section we list illustrative extensions to other problems of interest, starting with relatively straightforward steps and progressing towards exploratory possibilities.

Optimal Tuning of Algorithms. Many algorithms have tuning parameters. Optimizing these for a given fixed algorithm using similar formulations as presented is straightforward. Of course, alternative proposals exist, e.g., [18].

Rational Polynomial Algorithms. A rather obvious generalization is to consider not just a single monomial but rather rational polynomials involving the derivatives. More generally this could be extended to include non-integer and possibly even irrational powers. This would also allow algorithms involving Taylor-series expansion. No conceptual difficulty exists for an optimization formulation similar to the above but the number of variables increases.

Integral Operations. The formulation essentially also allows algorithms that perform integration, if $f^{(j)}$ with $j < 0$ is allowed. Obviously for multivariate programs ($n_x > 1$) the dimensionality needs to be appropriately considered.

Implicit Methods. Implicit methods, e.g., implicit Euler, do not only involve evaluation of derivatives but also the solution of a system of equations at each step. In other words we cannot directly write such methods as monomials in Definition 4. However, it is conceptually relatively easy to adjust the formulation, in particular by changing the update scheme to something along the lines of

$$x^{(it)} = x^{(it-1)} + \alpha \Pi_{j=0}^{j_{max}} \left(f^{(j)}(x^{(it)})\right)^{\nu_{j,1}}.$$

Computationally, the optimization problems will become substantially harder. The variables $x^{(it)}$ cannot be directly eliminated but rather the equations have to be given to the optimizer or addressed as implicit functions, see [21, 26].

Multistep Methods. Explicit multistep methods are those that calculate the current iterate based not only on the immediately previous but rather also on additional preceding steps. Examples include multistep integration methods as well as derivative-free equation solution methods such as

secant or bisection. It is straightforward to extend the formulations to allow such methods. Obviously the number of optimization variables increases accordingly. This can also be used for methods that rely on approximation of derivatives, e.g., Quasi-Newton methods. Similarly, methods such as secant can be encoded by allowing mixed polynomials of $\mathbf{x}$ along with $\mathbf{f}^{(j)}$. Bisection is not captured as this requires if-then-else; however, it should be relatively easy to capture this using integer variables which fits the proposed MINLP framework. Obviously for *implicit multistep* methods the two preceding formulations can be combined.

Global Optimization. A challenge in global optimization is that there are no explicit optimality criteria that can be used in the formulation. The definition of global optimality $f\left(\mathbf{x}^*\right) \leqslant f\left(\mathbf{x}\right)$, $\forall \mathbf{x} \in X$ can be added to the optimization of algorithm formulation but this results in a (generalized) semi-infinite problem (SIP). There are methods to deal with SIP problems but they are computationally very expensive [20]. Another challenge is that deterministic and stochastic global methods do not follow the simple update using just monomials as in Definition 4. As such, major conceptual difficulties are expected for such algorithms, including an estimation of the cost of calculating the relaxations.

Optimal Algorithms with Optimal Steps. Another interesting class are algorithms that include an optimization step within them, such as selection of an optimal step, e.g., [31, 33]. We can consider a bilevel formulation of the form in which the outer problem is similar to the formulation above, while the lower-level problem uses a different objective to calculate the optimal step. There exist (quite expensive) approaches to solve bilevel programs with nonconvexities [22]. With such a formulation we might be able to prove optimality of one of the three gradient methods described in [31] or discover alternatives. Additionally it would be possible to discover well-known methods such as the nonlinear conjugate gradient method [19], where the line search is described by the arg min operator or even stochastic methods which often depend on the computation of the maximum expected value of some iterative random variable $\mathcal{X}^{(it)}$. Similarly, it could be possible to find the desired controller algorithm in [12].

Continuous Form. Discrete methods are often reformulated in a continuous form. For instance, Boggs proposed a continuous form of Newton's method [7, 8] $\dot{x}(t) = \frac{f(x(t)}{f^{(1)}(x(t))}$. See also the recent embedding of discrete algorithms like Nesterov's scheme in continuous implementations [27, 30]. It seems possible to consider the optimization of these continuous variants of the algorithms using similar formulations. Typically discretization methods are used to optimize with such dynamics embedded. Thus, this continuous formulation seems more challenging to solve than the discrete form above. If a particular structure of the problem can be recognized, it could however be interesting, for instance to apply a relaxation similar to [25].

Dynamic Simulation Algorithms. The task here is to simulate a dynamical system, e.g., ordinary differential equations (ODE) $\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t))$ along with some initial conditions $\mathbf{x}(t=0) = \mathbf{x}^{init}$ for a given time interval $[0, t_f]$. We need to define what is meant by a feasible algorithm. A natural definition involves the difference of the computed time trajectory from the exact solution, which is not known and thus does not yield an explicit condition. One should therefore check the degree to which the resulting approximate curve, possibly after interpolation, satisfies the differential equation over this time interval.

Dynamic Optimization Algorithms. Dynamic optimization combines the difficulties of optimization and dynamic simulation. Moreover, it results in a somewhat amusing cyclic problem: we require an algorithm for the solution of dynamic optimization problems to select a dynamic optimization algorithm. As aforementioned, this is not prohibitive, e.g., in the offline design of an algorithm to be used online.

Algorithms in the Limit. Considering algorithms that provably converge to the correct solution in the limit makes the optimization problem more challenging. The infinite-dimension formulation (1) is in principle applicable with the aforementioned challenges. If the finite (parametrized) formulation (5), was directly applied, an infinite number of variables would have to be solved for. In such cases one could test for asymptotic self-similarity of the algorithm behavior as a way of assessing its asymptotic result.

Quantum Algorithms. A potential breakthrough in computational engineering would be realized by quantum computers. These will require new algorithms, and there are several scientists that are developing such algorithms. It would thus be of extreme interest to consider the extension of the proposed formulation to quantum algorithms and/or their real-world realizations. This may result in "regular" optimization problems or problems that need to be solved with quantum algorithms themselves.

Limitations

The proposed formulation is prototypical and a number of challenges are faced. As aforementioned, the knowledge of an explicit cost may not be a good approximation for all problems. It is also clear that different objective functions, including, e.g., considerations of memory usage or the inclusion of error correction features, may dramatically effect the results.

We expect the proposed optimization formulation to be very hard to solve and the numerical results confirm this. In particular we expect it to be at least as hard as the original problems. Proving that statement is outside the scope of the manuscript but two arguments are given, namely that the final constraint corresponds to solution of problem and

that the subproblems of the optimization problem are the same or at least related to the original problem.

Herein brute force and general-purpose solvers are used for the solution of the optimization problems. It is conceivable to develop specialized numerical algorithms that will perform much better than general-purpose ones due to the specific structure. In essence we have an integer problem with linear objective and a single nonlinear constraint. It seems promising to mask the intermediate variables from the optimizer. This would in essence follow the optimization with implicit functions embedded [21] and follow-up work. It is also conceivable to move to parallel computing, but suitable algorithms are not yet parallelized.

In our formulation, f is assumed to be a given function, so that we can apply current numerical methods for the optimization. It is, however, possible to consider multiple instances simultaneously, similar to stochastic optimization [6]. A simple method is to sample the instances of interest and optimize for some weighted/average performance of the algorithm. Alternatively the instances can be parametrized by continuous parameters and the objective in the above formulations replaced with some metric of the parameter distribution, e.g., the expectation of the cost. It is also possible, and likely promising, to consider worst-case performance, e.g., in a min-max formulation [13]. It is also interesting to consider the relation of this work to the theorems developed in [32], in particular that optimality of an algorithm over one class of problems does not guarantee optimality over another class.

Relation to Computational Complexity

As mentioned above, we expect our formulation to be at least as hard as the original problems. This would imply that the formulation cannot be easy to solve when the problems themselves are NP-hard.

Moreover, it is tempting to think that if the global solution of the formulation for an NP-hard problem across the family of all possible algorithms is an algorithm with exponential complexity, then we can conclude that P $\neq$ NP (if the best-possible algorithm is exponential, then no polynomial algorithm exists). However, it should be noted that we propose methods to solve for a fixed size NP-hard problem. The best algorithm for that size may exhibit exponential complexity, but this does not prove that there does not exist a polynomial algorithm which will be better for very large sizes. For instance, consider linear programs; simplex variants with exponential complexity outperform interior-point methods of polynomial complexity for all but the largest sizes. We could try to solve multiple instances of different sizes and explore how the performance grows. This would give some insight and good indications but most probably no proof (another algorithm could become better). If on the other hand we obtain a provably polynomial algorithm for a given NP-hard problem, this would imply that P=NP.

Conclusions

An MINLP formulation is proposed that can readily devise optimal algorithms (among a relatively large family of algorithms) for several prototypical problems, including solution of nonlinear equations and nonlinear optimization. Simple numerical case studies demonstrate that well-known algorithms can be identified along with new ones. We argue that the formulation is conceptually extensible to many interesting classes of problems, including quantum algorithms. Substantial work is now needed to develop and implement these extensions so as to numerically devise optimal algorithms for interesting classes of challenging problems where such algorithms are simply not known. The optimization problems obtained can, however, be very challenging and no claim is made that optimal algorithm discovery will be computationally cheap. However, in addition to the theoretical interest, there are certainly applications. Finding guaranteed optimal algorithms for a given problem implies understanding/classifying the difficulty of this problem. And it will certainly be worthwhile to automatically devise algorithms offline for problems to be solved online.

Acknowledgments. I.G.K. and A.M. are indebted to the late C.A. Floudas for bringing them together. Fruitful discussions with G.A. Kevrekidis and the help of C.W. Gear with the continuous formulation are greatly appreciated. The work of I.G.K. was partially supported by the US NSF.

Supporting Information (SI)

Solution of Equation Systems

For illustration purposes we will discuss the solution of equation systems using only zero and first derivatives. The zero derivative $f^{(0)}\left(\mathbf{x}^{(it)}\right)$ is a vector of the same size as the point $\mathbf{x}$. By definition, the required polynomial expression of the derivative $\left(f^{(0)}\left(\mathbf{x}^{(it)}\right)\right)^{\nu_k}$ is calculated element by element. Thus, we can equivalently rewrite each element if of the vector $\left(\mathbf{f}^{(0)}\left(\mathbf{x}^{(it-1)}\right)\right)^{\nu_k}$ as $\sum_{k=kmin}^{kmax} y_0^k \left(f_{if}^{(0)}\left(\mathbf{x}^{(it-1)}\right)\right)^k$ along with a constraint $\sum_{k=kmin}^{kmax} y_0^k = 1$. In contrast, the first derivative is the Jacobian matrix $\mathbf{J}$ with elements $\left.\frac{\partial f_{if}}{\partial x_{ix}}\right|_{\mathbf{x}^{(it)}}$. In principle the inverse can be written analytically but this is only realistic for small-scale systems. For systems of substantial sizes, it is better to calculate the direction via the solution of an equation system.

$$\mathbf{J}\left(\mathbf{x}^{(it-1)}\right)\Delta\mathbf{x} = \left(\mathbf{f}^{(0)}\left(\mathbf{x}^{(it-1)}\right)\right)^{\nu_k},$$

the solution of which gives us the direction $\Delta\mathbf{x}$ to be taken if the optimizer selects the inverse of the Jacobian, i.e., sets $\nu_1 = -1$. Similarly for other values of ν_1 we have to calculate other steps via explicit matrix multiplication or via similar equation systems. As aforementioned, the equations can be explicitly or implicitly written, see [21] and its extensions [26].

Alternative to Obtain a Finite Number of Variables

For the sake of simplicity an alternative is presented to obtain a finite number of variables in the optimal control problem. Again a maximal number of iterations it_{max} is considered and the cost is penalized by the residual, e.g.,

$$\min_{g,x^{(it)}} \sum_{it=1}^{it_{max}} \left|f\left(x^{(it)}\right)\right| \phi\left(g\left(x^{(it)}\right)\right)$$

$$\text{s.t. } x^{(it)} = g\left(x^{(it-1)}\right), \quad it = 1, 2, \ldots, it_{max} \tag{6}$$

$$\left|f\left(x^{(it_{max})}\right)\right| < \varepsilon \tag{7}$$

where again the last constraint ensures convergence. Without it, if algorithms with zero cost exist, then these will be chosen even if they are infeasible; they will give a zero objective value without meeting convergence.

Obtaining a Regular MINLP

The finite problem proposed contains terms $\left(f^{(j)}(x^{(it-1)})\right)^{\nu_j}$. We thus exponentiate a potentially negative real-valued $(f^{(j)}(x^{(it-1)}))$ to the power ν_j. At feasible points ν_j is integer and thus the operation well-defined. However, in intermediate iterations of the optimization, e.g., in branch-and-bound, ν_j may be real-valued and thus the operation is not defined. Consequently, the optimization formulation cannot be directly solved using standard MINLP solvers. It is however, relatively easy to reformulate as an MINLP with linear dependence on the integer variables. One way is to introduce for each j as many binary variables y_j^k as we have potential values for ν_j and enforce that exactly one is zero (special case of SOS-1 set). Then the term $f^{(j)}(x^{(it-1)})$ is equivalently rewritten as $\sum_{k=k_{min}}^{k_{max}} y_j^k \left(f^{(j)}(x^{(it-1)})\right)^k$ along with a constraint $\sum_{k=k_{min}}^{k_{max}} y_j^k = 1$.

Implementation

The implementation is performed in GAMS placing all the common elements placed in a main file. The problem-specific definitions are then given in corresponding GAMS files so it is easy to switch between problems. To limit the number of variables we avoid introducing variables for f and its derivatives and rather define them using macros. We do however introduce auxiliary variables for the iterates of the algorithm as well as variables that capture the residuals, see below. Both could be avoided by the use of further macros. Since we utilize global solvers we impose finite bounds for all variables.

For each iterate it we calculate the residuals by introducing a real-valued variable $res^{(it)}$ and a binary $y_{res}^{(it)}$ denoting if the desired convergence is met ($y_{res}^{(it)} = 0$) or not ($y_{res}^{(it)} = 1$). For the case of equation solving we take as residuals the value of the right-hand-side of the equation (function) $res^{it} = \max_{if} |f_{if}\left(\mathbf{x}^{(it)}\right)|$ and in the case of minimization the value of the derivative $res^{it} = \max_{ix} \left|\left(\frac{\partial f}{\partial x_{ix}}\right)\Big|_{\mathbf{x}^{(it)}}\right|$. This corresponds to the infinity norm and this is chosen since initial testing shows more robust optimization. The absolute values are written as pair of inequalities, e.g., for equation solving

$$res^{(it)} \geqslant \pm f_{if}\left(\mathbf{x}^{(it)}\right) \forall if.$$

Then the binary variable is decided based on a standard big-M constraint

$$y_{res}^{(it)} \geqslant (res^{(it)} - \varepsilon)/resscale$$

Only iterations that do not meet the residual tolerance are counted for in the objective

$$\min \sum_{it=1}^{it_{max}} y_{res}^{(it)} \left(\bar{\alpha}^{it} + \sum_{j=1}^{j_{max}} \sum_{k=k_{min}}^{k_{max}} y_j^k \phi_j^k\right)$$

For the cost constants ϕ_j^k we assume a zero cost for zero exponent, a cost of 1 for exponent 1, a cost of 1.5 for an exponent of 2, a cost of 2 for an exponent of -1 and a cost of 3 for an exponent of -2. The idea is to capture the expense of inversion and exponentiation.

These costs are multiplied by 1 for f, 10 for $f^{(1)}$ and 100 for $f^{(2)}$.

The problems we considered are

- $x^3 = 1$, $x \in [-2, 2]$, $x^{(0)} = 0.1$
- $x \exp(x) = 1$, $x \in [-2, 2]$, $x^{(0)} = 0.1$
- $\min_x x^4 + x^3 - x^2 - 1$, $x \in [-2, 2]$, $x^{(0)} = 0.1$
- Rosenbrock function in 2d $\min_{\mathbf{x}} 100(x_2 - x_1^2)^2 + (1 - x_1)^2$, $\mathbf{x} \in [-2, 2]^2$, $x_1^{(0)} = 0.7$, $x_2^{(0)} = 0.75$
- Simple 2d equation solving $x_2 - x_1^2 = 0$, $5x_2 - \exp(x_1) = 0$. Here we excluded second derivatives.

References

[1] S. Arya, D. M. Mount, N. Netanyahu, R. Silverman, and A. Y. Wu. An optimal algorithm for approximate nearest neighbor searching in fixed dimensions. In *Proc. 5th ACM-SIAM Sympos. Discrete Algorithms*, pages 573–582, 1994.

[2] R. Bacher. Automatic generation of optimization code based on symbolic non-linear domain formulation. In *Proceedings of the 1996 international symposium on Symbolic and algebraic computation*, pages 283–291. ACM, 1996.

[3] S. Bain, J. Thornton, and A. Sattar. Methods of automatic algorithm generation. In *PRICAI 2004: Trends in Artificial Intelligence*, pages 144–153. Springer, 2004.

[4] L. T. Biegler. *Nonlinear Programming: Concepts, Algorithms, and Applications to Chemical Processes*. MPS-SIAM Series on Optimization. SIAM-Society for Industrial and Applied Mathematics, 2010.

[5] P. Bientinesi. *Mechanical Derivation and Systematic Analysis of Correct Linear Algebra Algorithms*. PhD thesis, Graduate School of The University of Texas at Austin, 2006.

[6] J. R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer-Verlag, 1997.

[7] P. T. Boggs. The solution of nonlinear systems of equations by a-stable integration techniques. *SIAM J. Numer. Anal*, 8:767–785, 1971.

[8] P. T. Boggs. The convergence of the Ben-Israel iteration for nonlinear least squares problems. *Math. Comp.*, 30:512522, 1976.

[9] A. Brooke, D. Kendrick, and A. Meeraus. *GAMS: A User's Guide*. The Scientific Press, Redwood City, California, 1988.

[10] R. H. Byrd, J. Nocedal, and R. A. Waltz. *KNITRO: An Integrated Package for Nonlinear Optimization*, volume 83, pages 35–59. Springer-Verlag, 2006.

[11] C. P. Coelho, J. R. Phillips, and L. M. Silveira. Robust rational function approximation algorithm for model generation. In *Proceedings of the 36th annual ACM/IEEE Design Automation Conference*, pages 207–212. ACM, 1999.

[12] C. G. Economou. *An Operator Theory Approach to Nonlinear Controller Design*. PhD thesis, California Institute of Technology Pasadena, California, 1985.

[13] J. E. Falk and K. Hoffman. A nonconvex max-min problem. *Naval Research Logistics*, 24(3):441–450, 1977.

[14] B. A. Garber, D. Hoeflinger, X. Li, M. J. Garzaran, and D. Padua. Automatic generation of a parallel sorting algorithm. In *Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on*, pages 1–5. IEEE, 2008.

[15] J. R. Koza. *Genetic programming: on the programming of computers by means of natural selection*, volume 1. MIT press, 1992.

[16] J. R. Koza. Genetic programming ii: Automatic discovery of reusable subprograms. *Cambridge, MA, USA*, 1994.

[17] M. Kuhner, D. Burgoon, P. Keller, S. Rust, J. Schelhorn, L. Sinnott, G. Stark, K. Taylor, and P. Whitney. Automatic algorithm generation, Nov. 7 2002. US Patent App. 10/097,198.

[18] Q. Li, C. Tai, and E. Weinan. Dynamics of stochastic gradient algorithms. *arXiv:1511.06251*, 2015.

[19] D. G. Luenberger. *Introduction to linear and nonlinear programming*, volume 28. Addison-Wesley Reading, MA, 1973.

[20] A. Mitsos. Global optimization of semi-infinite programs via restriction of the right hand side. *Optimization*, 60(10-11):1291–1308, 2011.

[21] A. Mitsos, B. Chachuat, and P. I. Barton. McCormick-based relaxations of algorithms. *SIAM Journal on Optimization*, 20(2):573–601, 2009.

[22] A. Mitsos, P. Lemonidis, and P. I. Barton. Global solution of bilevel programs with a nonconvex inner program. *Journal of Global Optimization*, 42(4):475–513, 2008.

[23] L. Parada, M. Sepulveda, C. Herrera, and V. Parada. Automatic generation of algorithms for the binary knapsack problem. In *Evolutionary Computation (CEC), 2013 IEEE Congress on*, pages 3148–3152. IEEE, 2013.

[24] G. Ricart and A. K. Agrawala. An optimal algorithm for mutual exclusion in computer networks. *Communications of the ACM*, 24(1):9–17, 1981.

[25] S. Sager, H. G. Bock, and G. Reinelt. Direct methods with maximal lower bound for mixed-integer optimal control problems. *Mathematical Programming*, 118(1):109–149, 2009.

[26] M. D. Stuber, J. K. Scott, and P. I. Barton. Convex and concave relaxations of implicit functions. *Optimization Methods & Software*, 30(3):424–460, 2015.

[27] W. Su, S. Boyd, and E. J. Candes. A Differential Equation for Modeling Nesterov's Accelerated Gradient Method: Theory and Insights. *ArXiv e-prints*, Mar. 2015.

[28] M. Tawarmalani and N. V. Sahinidis. A polyhedral branch-and-cut approach to global optimization. *Mathematical Programming*, 103(2):225–249, 2005.

[29] D. C. Tozoni, P. J. de Rezende, and C. C. de Souza. A practical iterative algorithm for the art gallery problem using integer linear programming. 2013.

[30] A. Wibisono, A. C. Wilson, and M. I. Jordan. A Variational Perspective on Accelerated Methods in Optimization. *ArXiv e-prints*, Mar. 2016.

[31] A. Wibisono, A. C. Wilson, and M. I. Jordan. A variational perspective on accelerated methods in optimization. *arXiv preprint arXiv:1603.04245*, 2016.

[32] D. H. Wolpert and W. G. Macready. No free lunch theorems for optimization. *Evolutionary Computation, IEEE Transactions on*, 1(1):67–82, 1997.

[33] Y. Zhang, X. Chen, D. Zhou, and M. I. Jordan. Spectral methods meet em: A provably optimal algorithm for crowdsourcing. In *Advances in neural information processing systems*, pages 1260–1268, 2014.

©*A. Mitsos et al.*